Collaborative Expeditions through Remote Scientific Data via the Internet

Timothy A. Sandstrom and Velvin R. Watson, *Member, IEEE*

***Abstract* -- Current techniques for providing collaborative visual analysis of remote scientific data using the Internet are reviewed. It is shown why the tools in common use today, such as Microsoft's NetMeeting, do not provide scientists with highly interactive viewing of high resolution dynamic scenes of remote data when using typical Internet connections. New tools that do provide these capabilities in both synchronous and asynchronous collaborations have been developed and tested and are described in this paper.**

The significance of these new tools is the following capabilities they provide to scientists:

1. **The capability to access remote data and conduct individual or collaborative explorations through the data using highly interactive viewing of high resolution dynamic scenes.**
2. **The capability to record those explorations for later review by others.**
3. **The capability to publish reports on the Web containing a rich variety of guided expeditions (recorded explorations) through the remote data. Each expedition is automatically launched when it is selected from a Web page within their report.**
4. **The capability to experience other scientists' published expeditions through the data, the capability to modify or extend those expeditions with new "what if" explorations, and the capability to publish these new explorations back to the Web for others to experience.**
5. **The capability to experience an author's published expeditions through the data and then make a live connection with the author and/or other remote colleagues to conduct a synchronous collaborative extended exploration through the data. This new exploration can then be published to the Web.**

Testing of these tools proved that they provide highly interactive viewing of high resolution dynamic scenes of the data even in collaborations between continents.

***Index Terms* -- Internet, Collaborative Work, Remote Visualization, Scientific Visualization.**

I. Introduction

During the last decade, the Internet provided a major revolution in our ability to access remote information. The next major revolution being inspired by the Internet is in our ability to collaborate remotely.

Early efforts to enable collaboration using the Internet have focused mainly on video, audio, whiteboards, chat rooms, and document sharing tools. However, for a scientist, the tools that he/she commonly uses for analysis are often more important than any of these. Therefore, collaboration environments for scientists should also include the analysis tools commonly used by the scientist. Today, scientific visualization is a very popular analysis tool, so it is important to provide collaborative scientific visualization in the scientist's collaborative environment.

The reference list in this paper contains URLs pointing to a large number of research projects aimed at enabling collaboration over the Internet, and many of these projects are developing tools for collaborative scientific visualization. This paper describes the variety of techniques that are being applied to provide collaborative scientific visualization. The advantages and disadvantages of each are listed, and some guidelines for when to use each of the various techniques are given.

It is shown that the techniques used in the commonly available collaborative scientific visualization tools do not support highly interactive viewing of high resolution dynamic scenes of the data. This feature is very important for analysis of many types of data. Furthermore, with the high performance graphics now available even on PCs, this feature is now becoming expected in visual analysis tools. (For the remainder of this paper this feature, highly interactive viewing of high resolution dynamic scenes, will be referred to a "high performance visual analysis" and abbreviated as HPVA.)

Implementations of two techniques that do provide HPVA for collaborative scientific visualization are described in this paper. Results from testing these implementations are also presented.

These implementations make it relatively easy for scientists to conduct collaborative expeditions through remote scientific data.

II. Techniques being Applied for Remote Scientific Visualization

The references at the end of this paper provide a list of many current research projects in collaboration using the Internet. Many of these projects are developing tools for remote scientific visualization. The techniques employed in these projects differ mainly in what portion of the work in creating the pixel images is done remotely and what portion is done on the scientist's local computer.

The major steps in creating the pixel images are:

1. Read the data.
2. Convert the specific analysis of the data to scenes. For example, a scene might consist of isosurfaces of temperature. (The scenes are usually in the form of 3D scene graphs).
3. Convert the scene to 2D drawing primitives (usually triangles, lines, and dots).

4. Convert the drawing primitives to pixels to be shown on the monitor.

The techniques employed in the projects listed in the references are:

1. *Streaming Video from a Graphics Server* - Perform all of the steps listed above on a server attached to the remote data and send the pixels to the scientist's local computer for visualization. Usually the pixels files are compressed and sent as streaming video files with some loss of resolution from the original pixel images.
2. *Streaming Drawing Primitives from a Graphics Server* - Perform steps 1 through 3 listed above on a server attached to the remote data, send the drawing primitives to the scientist's local computer, and perform step 4 on the local computer. Systems employing X windows clients use this technique. Also, implementations of the proposed International Telecommunications Union (ITU) T.120 standard, such as Microsoft's NetMeeting, use this technique.
3. *Streaming Scene Graphs from a Graphics Server* - Perform steps 1 and 2 listed above on a server attached to the remote data, send the scene graphs to the scientist's local computer, and perform steps 3 and 4 on the local computer. Techniques based on the Virtual Reality Modeling Language (VRML)[30] and the 3D component of MPEG4[31] use this technique.
4. *Local Analysis of the Whole Data Set* - Send all of the data to the scientist's local computer and perform all of the steps above on the local computer.
5. *Local Analysis of Data Excerpts* - Send an excerpt of the data (such as only the data on a coarse grid over the region) to the scientist's local computer and perform all of the steps above on the local computer. Then send additional excerpts (such as the data on a fine grid over a small region) and analyze these data excerpts on the local computer. These subsequent excerpts are selected by using the analyses of earlier excerpts of the data.

For synchronized collaboration, it is desirable to be able to synchronize the views of all of the remote collaborators so they know exactly what their remote colleagues are looking at. For the *Streaming Video* and *Streaming Drawing Primitves* techniques, this can be accomplished by having the server send out the same video stream or the same drawing primitives to all of the scientists' local computers. For the *Streaming Scene Graphs* and the *Local Analysis* techniques, the application rendering the images on all of the scientists' local computers must be synchronized using application control commands sent over the Internet.

III. Advantages and Disadvantages of the Techniques

The advantages and disadvantages of these techniques are listed in the tables below.

TABLE I

ADVANTAGES AND DISADVANTAGES OF *STREAMING VIDEO FROM A GRAPHICS SERVER*

ADVANTAGES	DISADVANTAGES
The advantages are that the raw data isn't send over the Internet, the graphics power of the server can be used for the image rendering, and the scientists' clients only need to display video, so they can be dumb clients.	The major disadvantage is that interactive viewing of high resolution dynamic images requires a very high bandwidth between sites. Sending uncompressed images with 24 bits of color per pixel and 1280x1024 pixel resolution at 30 frames per second would require a bandwidth of 1 Gbps. Since most scientists have far less than 1 Gbps Internet bandwidth available to them and are afraid of artifacts from highly compressed images, they will have to sacrifice interactive viewing of dynamic scenes or high resolution. In other words, tools developed with this technique will not provide HPVA to the scientists with typical Internet bandwidths.

TABLE II

ADVANTAGES AND DISADVANTAGES OF *STREAMING DRAWING PRIMITIVES FROM A GRAPHICS SERVER*

ADVANTAGES	DISADVANTAGES
The advantages are that the raw data isn't send over the Internet, the graphics power of the server can be used to create the graphics primitives, and the scientists' clients only need to have good 2D graphics capability --- high performance 3D graphics capability is not required.	The major disadvantage is that complex scenes require a large number of drawing primitives, and in some cases the size of the drawing primitives file can exceed the size of the pixel file generated from the drawing primitives. Therefore, interactive viewing of complex, high resolution dynamic scenes for this technique requires a very high bandwidth between sites. Scientists with typical Internet bandwidths must sacrifice interactive viewing of dynamic scenes, scene complexity, or high resolution. In other words, tools developed with this technique will not provide HPVA to the scientists with typical Internet bandwidths.

TABLE III

ADVANTAGES AND DISADVANTAGES OF *STREAMING SCENE GRAPHS FROM A GRAPHICS* SERVER

ADVANTAGES	DISADVANTAGES
The advantages are that the raw data isn't send over the Internet, the graphics power of the server can be used to create the scene graphs, and the scientist does not have to sacrifice highly interactive viewing or high resolution *if the scene graphs are not dynamic*. (If the scene graphs are dynamic, then the scientist may have to sacrifice interactive viewing of the dynamic scenes, scene complexity, or high resolution.)	The major disadvantage is that complex scenes also require large scene graph files. However, these files are usually much smaller than either the pixel files or the drawing primitive files. If the scenes are complex and rapidly changing, this technique also requires a high bandwidth. Therefore, scientists with typical Internet bandwidths may have to sacrifice interactive viewing of dynamic scenes, scene complexity, or high resolution. In other words, tools developed with this technique will not provide HPVA to the scientists with typical Internet bandwidths. Another disadvantage is that the scientist's local computer must have good 3D graphics rendering capabilities if the scenes are 3D. (Fortunately, good 3D graphics capabilities are now becoming available even on PCs.)

TABLE IV

ADVANTAGES AND DISADVANTAGES OF *LOCAL ANALYSIS OF THE WHOLE DATA SET*

ADVANTAGES	DISADVANTAGES
The major advantage is that the only restriction to HPVA is the power of the graphics card on the scientist's local computer. Fortunately, new graphics cards can now provide very effective HPVA even on PCs. HPVA is typically much more effective than the analyses that can be achieved by the three streaming techniques listed in Tables I through III above.	The major disadvantage is that the raw data must be sent over the Internet and stored on the scientist's local computer. However, this is usually done prior to the analysis phase and does not impact the ability for HPVA. Another disadvantage is that the scientist's local computer must have good 3D graphics rendering capabilities if the scenes are 3D. (Fortunately, good 3D graphics capabilities are now becoming available even on PCs.)

TABLE V

ADVANTAGES AND DISADVANTAGES OF *LOCAL ANALYSIS OF DATA EXCERPTS*

ADVANTAGES	DISADVANTAGES
When analyzing each data segment, the major advantage is the same as the *Local Analysis of the Whole Data Set* technique listed above plus the advantage that not all of the data must fit on the scientist's local computer.	The major disadvantage is that excerpts of the raw data must be sent over the Internet and stored on the scientist's local computer. In this case, not all of the data is sent before the analysis phase, so there can be a substantial delay in the analysis phase each time a new data excerpt is requested. However, for each data segment, this does not impact the ability for HPVA. Another disadvantage is that the scientist's local computer must have good 3D graphics rendering capabilities if the scenes are 3D. (Fortunately, good 3D graphics capabilities are now becoming available even on PCs.)

IV. GUIDELINES FOR WHEN TO USE EACH TECHNIQUE

Following are the key factors for deciding between the techniques above.

1. The power of the scientist's local computer (Dumb local computers vs powerful local computers).
2. The power and responsiveness of the remote computer (Overloaded, non-responsive servers vs powerful, responsive servers).
3. The types of scenes to be rendered in the analysis (Scenes with typical rendering provided by standard graphics cards vs scenes with more exotic rendering, such as many bounce ray tracing).
4. The degree of interactivity required (Non-interactive visualizations vs highly interactive visualizations).
5. The dynamics of the scenes to be viewed (Static vs highly dynamic scenes).
6. The dimensionality of the scenes to be viewed (2D vs 3D scenes).
7. The size of the remote data files (Large data files vs small data files).

When applying these factors, one should recognize that the importance of HPVA is often not appreciated by scientists during the initial phase of an analysis. Frequently just a static 2D plot is all that is requested. However, once the scientists learn to take advantage of HPVA, they appreciate the advantages it offers and often begin to specify it as a requirement in their future visualization tools.

If the scientist's local computer is only an X windows device or similar dumb terminal, then one must use either the *Streaming Video* or *Streaming Drawing Primitives* technique. However, most scientists now have at least a reasonably powerful PC, so for the rest of this discussion we will assume that a reasonably powerful local computer is available.

If the remote server for processing the data is anemic or overloaded and non-responsive, then one should use either of the *Local Analysis* techniques. However, for the rest of this discussion we will assume a reasonably powerful and responsive remote server is available for remote processing.

If the scenes must be rendered with techniques not available on the graphics card in the scientist's local computer, then the *Streaming Video* technique using a powerful graphics server is likely to be required. For the rest of this discussion, we will assume that the scenes can be rendered with the gourand shaded triangle rendering available on most cards today. If the scenes are not more complex than those in the current computer games, then PC graphics cards can probably render them quickly. If volume rendering is required, there are also special PC cards now available that can handle that.

If the visualizations are non-interactive, then it is probably best to create static scenes or non-interactive movies using a standard image transfer or the *Streaming Video* technique. The local computer can then be used to just flip through the cached static scenes or non-interactive movies.

For highly interactive visualizations of 2D static scenes, it is best to use the *Streaming Drawing Primitves* or *Streaming Scene Graphs* techniques. These are better than the *Streaming Video* technique because one can zoom or pan through the scene using the local computer to render the zoom or pan windows without retransmitting any information from the remote computer.

For highly interactive visualizations of static 3D scenes, it is best to use the *Streaming Scene Graph* technique. This permits interactive manipulation of the viewing position using the local 3D rendering capability without retransmitting any information from the remote computer. Viewers for VRML and the 3D component of MPEG4 use this technique.

For HPVA, whether 2D or 3D, it is best to use the *Local Analysis of the Whole Data Set* technique if the data will fit on the local computer or the *Local Analysis of Data Excepts* technique if the data will not fit on the computer. That is because all the streaming techniques cannot provide HPVA with the Internet bandwidths that are typically available to scientists today.

V. IMPLEMENTATIONS OF THE STREAMING TECHNIQUES

These implementations are common and are only described briefly here because none of them provide collaborative

HPVA with Internet bandwidths that are typically available. This is a major disadvantage for analysis of many types of data.

The most common implementation of the *Streaming Video* technique is the use of Web tools for publishing non-interactive movies on the Web for asynchronous collaboration

The most common implementations of the *Streaming Drawing Primitives* technique are the use of X windows and the use of the proposed ITU T.120 standard for both synchronous and asynchronous collaborations. Examples of the ITU T.120 standard implementations are the popular Microsoft NetMeeting, Sun's SunForum, SGI's SGIMeeting, and HP's HP Visual Conference. These implementations have the great advantage that they can be used to collaborate across platforms and that they are prevalent --- especially NetMeeting.

The most common implementations of the *Streaming Scene Graphs* technique are the VRML style viewers, including the 3D component of the MPEG4 standard. Although the MPEG4 viewers are not yet prevalent, it is likely that they will become very common within the next year.

VI. IMPLEMENTATIONS OF THE *LOCAL ANALYSIS OF THE WHOLE DATA SET* TECHNIQUE

The importance of implementations of this technique is that they offer the greatest potential for collaborative HPVA for analysis of data that will fit on the scientist's local computer.

A. RemoteFAST and FASTexpeditions

This technique was implemented by NASA Ames Research Center as RemoteFAST for synchronous collaborations and as FASTexpeditions for asynchronous collaborations in computational physics. These are extensions of FAST[21] (Flow Analysis Software Toolkit), which was also developed at NASA Ames Research Center. The design and testing of these tools is described in detail in reference [34].

The implementations for both the synchronous and asynchronous collaboration tools utilized FAST's journal file capability, wherein a script of each exploration through the data is automatically recorded in a journal file as the scientist conducts the exploration. These scripts can then be used to replay the exploration. These recorded explorations are called expeditions. The scripts are just ASCII files so they can be easily edited by an editor.

1) Implementation for asynchronous collaboration

To create an asynchronous collaborative visualization tool, FAST was wrapped with a C Shell script to permit use with the World Wide Web. This tool is named FASTexpeditions. The data to be analyzed and the expeditions (journal files) are made available from Web pages. Selecting the data from a Web page causes downloading of the data to the local computer, automatic launching of FASTexpeditions on the local computer, and execution of a script to set up the initial state of the analysis. Subsequent selections of expeditions from the Web page causes execution of the journal files for those recorded explorations.

For most of the investigations that we have posted to the Web, all of the expeditions (journal files) are packaged and downloaded along with the initial data because doing this permits playing of any expedition without returning to the remote Web server for the journal files. In this case, the URL used on the Web page refers to the downloaded journal files on the local disk, so the Web browser gets these immediately from the local computer disk rather than waiting for the remote Web server to respond and deliver them.

Sound files can be included in the expeditions for an audio description of the analysis as it occurs.

To provide safety from people who might post malicious journal files, the C Shell wrapper scans each journal file and removes unsafe commands.

To facilitate the ease of collaboratively discussing the posted analyses with remote colleagues, the Web pages containing the FASTexpeditions also contain selections for automatically initiating a synchronous collaboration using RemoteFAST (described in the next section).

A utility was created to automatically generate a FASTexpedition Web page with URLs pointing to the data from the computer simulation and the journal files of the expeditions.

2) Implementation for synchronous collaboration

To create a synchronous collaborative visualization tool, FAST was combined with a program to handle TCP/IP unicast peer to peer communications between remote sites. This tool was named RemoteFAST [22]. (As soon as multicast is prevalent, it should be used instead of unicast for multipoint collaborations to eliminate the need to send multiple streams from the controlling site.) To start a synchronous session, the data to be analyzed is distributed to each site and FAST is launched at each site. (This is usually done automatically by using FASTexpeditions, described in the previous section.) Then, the program to handle communications at each site is launched to create a daemon dedicated to efficient passing of events between the sites. During the session, the controlling RemoteFAST site simply detects the script commands as they are being recorded into the journal file and sends the same script commands over the network to all controlled RemoteFAST sites. At the controlled sites, RemoteFAST simply reads the incoming script commands as though they were being read from recorded journal files on the local disk and passes them onto FAST.

This implementation provides many advantages. It is simple. The bandwidth between sites need not be large because only script commands are sent between sites. And, the system response experienced by the users is nearly the same as the response in stand-alone mode. The system response is very good because:

1. The dedicated communications daemons provide a nearly unnoticeable delay in sending the script commands over the network.
2. Intelligent, compact information (i.e., application specific data and events rather than pixels) is sent between sites.

3. The 3D scene rendering is performed by the local computer.

Therefore, all remote scientists appear to be seeing the same high resolution, dynamic, 3D scenes simultaneously.

RemoteFAST is normally used along with a desktop video tool if the network bandwidth permits, or along with a normal phone conference if the network bandwidth doesn't permit the video.

These remote collaboration sessions can be recorded and posted onto the Web for other scientists to playback and modify at their convenience.

3) Advantages of FASTexpeditions and RemoteFAST

The advantages of these new tools are the following capabilities they provide to scientists:

1. The capability to access remote data and conduct individual or collaborative explorations through the data using HPVA.
2. The capability to record those explorations for later review by others.
3. The capability to publish in their scientific reports on the Web a rich variety of guided expeditions (recorded explorations) through the remote data. Each expedition is automatically launched when it is selected from a Web page within their publication.
4. The capability to experience other scientists' published expeditions through the data, the capability to modify or extend those expeditions with new "what if" explorations, and the capability to publish these new explorations back to the Web for others to experience.
5. The capability to experience an author's published expeditions through the data and then make a live connection with the author and/or other colleagues to conduct a synchronous collaborative extended exploration through the data. This new exploration can then be published to the Web.

These tools are available for download at no cost from http://www.nas.nasa.gov/Software/FAST/FASTexpeditions.

4) Disadvantages of FASTexpeditions and RemoteFAST

In addition to the disadvantages inherent in the *Local Analysis of the Whole Data Set* technique, these tools also have the disadvantage that they require an SGI workstation. Although FAST and SGI workstations have many users in the computational physics field, they are not prevalent in some other fields.

B. Gel

Gel is a new NASA visualization tool developed for data analysis that does not require SGI workstations. This package provides an environment for doing research in visualization techniques such as scalar and vector field detection, topological analysis, and surface flow patterns. In addition, it allows the NASA Ames Data Anlysis Group to research various methodologies such as time varying visualization, direct manipulation, distributed computing, remote collaboration, out-of-core visualization, and multithreaded library design. Gel runs on SGI, Linux, and Windows NT systems.

1) Implementation for asynchronous collaboration

Gel's user interface, similar to that of FAST, creates a journal file of all user interactions. This ASCII text file could be executed at a later time to recreate the results of a session. Thus by sharing the scripts as well as the data, asynchronous collaboration is supported much in the same way as FAST. Using an intranet web portal, one could allow trusted users access to shared data and scripts.

2) Implementation for synchronous collaboration.

Gel allows a user to connect his/her session to another running on a separate machine using a TCP/IP unicast socket interface. When two sessions are connected this way in a peer-to-peer relationship, the entries sent to the journal file are also sent via a socket interface to all other peers in the group. At the time of this writing, each peer must establish a connection manually to all other peers. This should be enhanced to enable a mode whereby connecting to one of the peers automatically connects you to all other peers in the session. Currently, no security methods are in place to provide authentication for group membership. Nor is there a method in place for arbitration when users execute conflicting commands such as if one person were to delete an object while another is changing it's execution parameters; this is an area of ongoing research.

VII IMPLEMENTATIONS OF THE *LOCAL ANALYSIS OF DATA EXCERPTS* TECHNIQUE

The importance of implementations of this technique is that they offer the potential for HPVA for analysis of data that will not fit on the scientist's local computer.

Gel, described in the previous section, has been modified to support automatic paging of the remote data. These methods build on the capabilities of remoteFAST and FASTexpeditions in two ways. First, star topologies are possible whereby multiple users can interact with the same 3D scene at once. Secondly, with the capabilities of a new method of demand paging [32] from a remote machine, collaborators have the potential for viewing data that resides in a location remote to all the viewers. This would allow for instance, the data to reside on a central server with large disk capacity. Granted, this technique's viability would greatly depend upon the capabilities of both the processing power and the network speed of the data server. But in some exceptional cases, a high speed server with a large, very fast disk array is capable of delivering data over a LAN to a client machine faster than the client is able to read the same data off it's own, albeit slower, disk drives [33]. Demand Paging is similar to virtual memory systems where data is only brought into memory when it is touched by the application. Pages read into memory (and in this case, over a network) are cached, and replaced under a Least Recently Used schema. This technique leverages the fact that some visualization techniques common to Computational Fluid Dynamics such as particle tracing touch only a small fraction of the total data set. Also, Demand Paging allows the user to work with data that would ordinarily be too big to view on their local machine. We have performed interactive visualization of 5 to 10 GB data sets on systems with 500 MB to 1 GB of memory [33].

VIII RESULTS FROM TESTING THE TWO IMPLEMENTATIONS

A. RemoteFAST and FASTexpeditions

RemoteFAST and FASTexpeditions have been tested in collaborative sessions between sites within the U.S. and between sites in different continents Within the U.S., the tests were conducted primarily between the NASA Ames Research Center in California and the EPA (Environmental Protection Agency) in North Carolina. Tests between the U.S. and Australia were conducted between NASA Ames Research Center and Perth Australia. Tests between the U.S. and Europe were conducted between the EPA and Monte Carlo, Monaco or Poitiers, France. Figure 1 shows the computer screen during a session.

RemoteFAST and FASTexpeditions were highly effective for both synchronous and asynchronous collaboration. The effectiveness of the collaboration was nearly as good as being together in the same office and looking at the same workstation while using FAST for the analysis or for a playback of an analysis.

For synchronous collaboration, the response of the visual analysis tool was nearly the same as in stand-alone mode. All sites were able to view the same high resolution (1280x1024), dynamic, 3D scenes simultaneously. Individual sites could independently control their own scene viewing position, but the viewing position could also be resynchronized with the controlling site's viewing position. Control of the analysis was easily transferred between sites. The bandwidth utilized between sites during a remote collaboration session was measured to peak at less than 1K bit/second. Note that this low bandwidth utilization and high display performance is achieved by sending script commands over the network and by having the local computer create and render the scenes. This performance cannot now be achieved by sending pixels over the network. Even systems that send scene graphs (such as VRML files) over the network do not match this performance.

For asynchronous collaboration, the analyses posted on the Web were easily downloaded and played. After the initial data download, the playback performance was identical to the performance of playback from journal files on the local disk.

Stereo glasses were often used to obtain stereoscopic scenes in both synchronous and asynchronous modes.

The major advantages of FASTexpeditions over VRML or movie files posted on the Web are:

1. The 3D display performance is superior.
2. Viewers download the actual data and can perform their own "what if" analysis on the data.
3. Viewers can modify the analyses they download and post their own analyses back on the Web.
4. Viewers can collaboratively review and modify the posted analyses with remote colleagues, and these analyses can be posted back onto the Web.

RemoteFAST and FASTexpeditions were used in conjunction with InPerson™, SGI's desktop video conference tool, whenever the network bandwidth was high enough (i.e., between France and the U.S. and between sites within the U.S.). Ordinary phones were used instead of InPerson™ when the network bandwidths would not support satisfactory desktop video (i.e., between Monaco and the U.S. and between Australia and the U.S.).

The scenario used most often to demonstrate the features of FASTexpeditions and RemoteFAST follows:

1. A scientist goes to a Web site where FASTexpeditions of various analyses of computer simulations of physics are posted.
2. The scientist selects one of the FASTexpeditions and views several of the posted analyses of the data.
3. The scientist then extends the author's posted analysis with his/her own "what if" analysis.
4. The scientist then contacts the author of the posted analyses with a phone or InPerson™ and asks the author about one of the features seen in an analysis.
5. The author and the scientist then both initiate a remote collaboration by making selections on the Web page to automatically start RemoteFAST.
6. The author and the scientist then use RemoteFAST collaboratively to investigate the feature.

Typically, the desktop video was only used at the beginning of the collaborative session when establishing initial contact. When the interest shifted from the initial "hello" to the analysis of the data, the primary focus was shifted to the 3D scenes of the visual analysis process and to the audio.

B. Gel

Gel has been tested on the NAS LAN. In these tests, Gel provided highly interactive viewing of high resolution dynamic scenes. The paging of data from the high speed remote server was approximately the same speed as from a local disk on the workstation. Tests on a variety of remote network conditions are still in progress.

IX CONCLUSIONS

Collaboration over the Internet has the potential for making a major impact on the way we conduct scientific research. To achieve this potential, it is important to include the scientists' analysis tools within their collaborative environment. The tools described in this paper provide collaborative visual analysis of remote scientific data using the Internet.

These tools have been tested and found to provide highly interactive viewing of high resolution dynamic scenes even in collaborations between continents.

ACKNOWLEDGMENTS

FAST was created at NASA Ames Research Center by the team listed in [21]. Jean Clucas and John West were especially helpful in creating FASTexpeditions and

RemoteFAST[22]. Todd Plessel, with the EPA in Research Triangle Park in North Carolina, installed these programs within the EPA and helped extensively with testing and with many demonstrations.

Gel was created by the lead author using components created by the NAS Data Analysis Group. David Ellsworth, creator of the remote-paging technique, implemented the remote paging server and advised on performance and architectural issues.

REFERENCES:

Current Research on Collaborative Problem Solving Environments by Organization

Department of Energy

[1] Common Component Architecture (CCA)
(CCA provides components especially valuable for scientific research)
Contact - Rob Armstrong - Sandia
Website - http://z.ca.sandia.gov/~cca-forum/

[2] Collaboration Technologies Group
(General research and development in collaborat ion technologies)
Contact - Deb Agarwal - LBL
Website - http://www-itg.lbl.gov/Collaboratories/

[3] Collaboratory Interoperability Framework Project (CIF)
(CIF will provide fundamental communications needs to build on)
Contact - Deb Agarwal - LBL
Website - http://www-itg.lbl.gov/CIF/

[4] Advanced Visualization Communications Toolkit
(This will provide network aware components to optimize visualization)
Contact - Deb Agrawal - LBL
Website-http://www-itg.lbl.gov/~deba/NGI/AdvVizComm.html

[5] Corridor One Project
(High performance visualization over very high bandwidth networks)
Contact - Rick Stevens - ANL
Websitehttp://www-fp.mcs.anl.gov/fl/research/Proposals/co.htm

[6] Toolkit for Collaboratory Development
(Includes Core2000 – a collaboratory research environment based on Habanero from NCSA)
Contact – Jim Meyers – PNL
Website - http://www.emsl.pnl.gov:2080/docs/collab/

[7] DOE2000 Electronic Notebook Project
(Electronic notebook projects at PNL, LBL, and ORNL)
Main Website - http://www.epm.ornl.gov/enote/
Contacts
Jim Myers - PNL
Sonia Sachs- LBL
Al Geist – ORNL

[8] DOE2000 Collaboratory Research
(Basic collaboratory research projects)
Website - http://www-unix.mcs.anl.gov/DOE2000/collabs.html

[9] DOE2000 Collaboratory Pilot Projects
The Diesel Combustion Collaboratory
Website - http://www-collab.ca.sandia.gov/dcc/
The Materials MicroCharacterization Collaboratory
Website - http://tpm.amc.anl.gov/MMC/
Environmental Molecular Sciences Collaboratory
Website - http://www.emsl.pnl.gov:2080/docs/collab/
Fusion Collaboratory
Website - http://www.fusionscience.org/collab/REE/

NSF and Universities

[10] WebFlow
(A visual programming paradigm for Web/Java based coarse grained distributed computing. This is based on Java technology.)
Contacts
Tomasz Haupt - Syracuse University
Wojtek Furmanski - Syracuse University
Website - http://www.npac.syr.edu/users/haupt/WebFlow/

[11] SCIRun
(This provides computational steering)
Contact - Christopher R. Johnson - University of Utah
Website - http://www.sci.utah.edu/publications/scitools96/node1.html

[12] NCSA Collaboration Systems
(Habanero provides state and event synchronization for multiple copies of a software tool. It utilizes Java.)
Contact - Polly Baker - NCSA
Website - http://havefun.ncsa.uiuc.edu/

[13] PUNCH
(An Architecture for Web-Enabled Wide-Area Network -Computing)
Contact - Prof Jose Fortes - Purdue Univ
Website - http://punch.ecn.purdue.edu/

[14] DISCIPLE
(Distributed System for Collaborative Information Processing and Learning)
Contact - Dr. Ivan Marsic - Rutgers Univ
Website - http://www.caip.rutgers.edu/multimedia/groupware/

[15] Space Physics and Aeronomy Research Collaboratory
(an environment for collaborative research in space physics and aeronomy)
Contact - Gary Olson - Univ of Michigan
Websit e - http://intel.si.umich.edu/sparc/

[16] Stanford Interactive Workspaces
(For exploring possibilities for people to work together in technology-rich spaces)
Contact - Terry Winograd - Stanford
Website - http://graphics.stanford.edu/projects/iwork

[17] High Performance and Real Time Corba
(Research on improving the throughput perfomance and reducing latency of Corba)
Contact - Doug Schmidt - UC Irvine (Wash Univ.)
Website - http://www.cs.wustl.edu/~schmidt/corba-research-performance.html

NASA

[18] Intelligent Synthesis Environment (ISE) and Collaborative Engineering Environment (CEE)
(NASA's project to create collaborative analysis and design environments)
Contact – W Lundy, NASA Lewis Research Center, for ISE
Website - http://ise.nasa.gov
Contact - Ed Chow, NASA Jet Propulsion Laboratory, for CEE
Website - http://ce-server.jpl.nasa.gov/

[19] Science Desk
(A project to create collaborative research environments with AI support)
Contact - Rich Keller – NASA Ames Research Center
Website - http://sciencedesk.arc.nasa.gov

[20] Mars Web Pages
(A website for the collaborative selection of Mars landing sites)
Contact - Glenn Deardorff – NASA Ames Research Center
Website – http://marsoweb.nas.nasa.gov/landingsites/

[21] FAST (Flow Analysis Software Toolkit
(A tool for visual analysis of computer simulations of complex physics)
Contact – Tim Sandstrom, NASA Ames Research Center
Website – http://www.nas.nasa.gov/Software/FAST

[22] RemoteFAST and FASTexpeditions
(Tools for asynchronous and synchronous collaborative scientific visualization)
Contact – Val Watson – NASA Ames Research Center
Website http://www.nas.nasa.gov/Software/FAST/FASTexpeditions

Industry

[23] Intelligent Human-Computer Interaction
(An environment for collaboration based on rooms. Awareness and privacy issues are addressed.)
Contact - Samuel Bayer - Mitre Corp
Website - http://www.mitre.org/resources/centers/it/g063/hci-index.html

Commercial CPSE Systems

[24] Tango Interactive
(Based on WebFlow from Syracuse University)
Contact - Marek Podgorny - WebWisdom
Website - http://www.webwisdom.com/tangointeractive/

CPSE Organizations

[25] Computingportals
Home Page - http://www.computingportals.org/
Survey of projects - http://www.computingportals.org/projects

[26] Computer Supported Cooperative Work (CSCW)
Website - http://www.acm.org/sigchi/cscw2000/index.html

Reports on CPSEs

[27] Report on Collaborative Virtual Environments 1998
University of Manchester, UK, 17-19th June 1998
Elizabeth Churchill and David Snowden
http://www.fxpal.xerox.com/ConferencesWorkshops/cve/Report.htm

[28] Workshop on CPSEs for Scientific Research
San Diego, CA., 29 June – 1 July, 1999
http://www.emsl.pnl.gov:2080/docs/cpse/workshop/index.html

CPSE Standards

[29] International Telecommunication Union's Proposed Standards
Complete listing of proposed standards -
http://www.itu.int/publications/telecom.htm
Proposed standard for application sharing –
http://www.itu.int/itudoc/itu-t/rec/t/t120.html
Proposed standard for audiovisual and multimedia systems -
http://www.itu.int/itudoc/itu-t/rec/h/h323.html

3D Web Standards

VRML

[30]VRML Website – http://www.web3dorg/

MPEG4

[31] MPEG4 Website – http://www.ccir.ed.ac.uk/mpeg4/

Other References

[32] M. B. Cox and D. Ellsworth. Application-controlled demand paging for out-of-core visualization. In Roni Yagel and Hans Hagen, editors, IEEE Visualization '97, pages 235 - 244. IEEE, October 1997

[33] D. Ellsworth. Accelerating Demand Paging for Local and Remote Out-of-Core Visualization. Submitted to IEEE Visualization '01. Currently In Review.

[34] V. R. Watson, "Supporting Scientific Analysis within Collaborative Poblem Solving Environments", HICSS-34 Minitrack on Collaborative Problem Solving Environments, Maui, Hawaii, January 3-6, 2001.
http://www. nas.nasa.gov/~watson/HICSS34paper.html